

CLAIMS

What is claimed is:

- 1 1. A multi-node computer system comprising:
 - 2 a local node, the local node comprising at least one local sub-node, the at least one local
 - 3 sub-node including a first local sub-node, the first local sub-node comprising:
 - 4 a local dynamic memory, the local dynamic memory being a volatile system
 - 5 memory used by a processor in the first local sub-node;
 - 6 a scalability port including a write-through transmit buffer associated with the
 - 7 local dynamic memory; and
 - 8 a first scalability chipset comprising a first memory controller that directs a write
 - 9 of data to the local dynamic memory simultaneous with a back-up write of the
 - 10 data to the write-through transmitting buffer;
 - 11 a remote node, the remote node comprising at least one remote sub-node, the at least one
 - 12 remote sub-node including a first remote sub-node, the first remote sub-node comprising:
 - 13 a back-up memory for the local dynamic memory of the first local sub-node, the
 - 14 back-up memory being distinct from local system memory of the first remote sub-
 - 15 node;
 - 16 a receiving interface buffer for receiving, from the write-through transmit buffer
 - 17 associated with the local dynamic memory of the first sub-node, data written to
 - 18 the local dynamic memory of the first sub-node; and
 - 19 a second scalability chipset comprising a second memory controller that directs a
 - 20 write to the back-up memory of data received at the receiving interface buffer;
 - 21 and
 - 22 an input/output controller in the second scalability chipset that assigns a location identity
 - 23 of the first local sub-node to a replacement sub-node only if the first local sub-node is
 - 24 removed from the multi-node computer system, the location memory based on a memory
 - 25 map included in the data stored in the back-up memory in the first remote sub-node.
- 1 2. The multi-node computer system of claim 1, wherein the replacement sub-node is the
- 2 first remote sub-node.

1 3. The multi-node computer system of claim 1, wherein the replacement sub-node is a
2 second remote sub-node of the remote node, the second remote sub-node having received and
3 stored the data from the back-up memory of the first remote sub-node.

1 4. The multi-node computer system of claim 1, wherein the replacement sub-node is a
2 second local sub-node of the local node, the second local sub-node having received and stored
3 the data from the back-up memory of the first remote sub-node.

1 5. The multi-node computer system of claim 1, wherein the first local sub-node is removed
2 as a hot swap, in which all other nodes of the multi-node computer system remain non-quiescent.

1 6. A method of removing a node from a multi-node computer, the method comprising:
2 receiving a system management interrupt (SMI) in a node in a multi-node computer;
3 quiescenting only the node receiving the SMI;
4 polling other nodes in the multi-node computer to determine if the SMI affects an
5 operation of any of the other nodes;
6 quiescenting any other SMI affected node; and
7 transferring all of the contents of any affected node's system memory to a backup
8 memory in an unaffected node in the multi-node computer, wherein the unaffected node assumes
9 all operations of the node that received the SMI, thus allowing the node to be removed from the
10 multi-node computer.

1 7. The method of claim 6, wherein the SMI is in response to a request to hot-swap out the
2 node.

1 8. The method of claim 6, wherein the SMI is in response to a predicted failure of the node.

1 9. A method of removing a sub-node from a multi-node computer, the method comprising:
2 Receiving a system management interrupt (SMI) in a sub-node in a multi-node computer;

3 quiescenting only the sub-node receiving the SMI;
4 polling other sub-nodes in the multi-node computer to determine if the SMI affects an
5 operation of any of the other sub-nodes;
6 quiescenting any other SMI affected sub-node; and
7 transferring all of the contents of each affected sub-node's system memory to a backup
8 memory in an a respective unaffected sub-node in the multi-node computer, wherein the
9 unaffected sub-node assumes all operations of the sub-node that received the SMI, thus allowing
10 the sub-node to be removed from the multi-node computer.

1 10. The method of claim 9, wherein the unaffected sub-node is in a node that does not
2 include the SMI affected sub-node.

1 11. The method of claim 9, wherein the unaffected sub-node is in a same node as the SMI
2 affected sub-node.

1 12. A computer program product, residing on a computer usable medium, for removing a
2 node from a multi-node computer, the computer program product comprising:
3 program code for receiving a system management interrupt (SMI) in a node in a multi-
4 node computer;
5 program code for quiescenting only the node receiving the SMI;
6 program code for polling other nodes in the multi-node computer to determine if the SMI
7 affects an operation of any of the other nodes;
8 program code for quiescenting any other SMI affected node; and
9 program code for transferring all of the contents of any affected node's system memory to
10 a backup memory in an unaffected node in the multi-node computer, wherein the unaffected
11 node assumes all operations of the node that received the SMI, thus allowing the node to be
12 removed from the multi-node computer.

1 13. The computer program product of claim 12, wherein the SMI is in response to a request
2 to hot-swap out the node.

- 1 14. The computer program product of claim 12, wherein the SMI is in response to a
2 predicted failure of the node.